

Event-Driven Analysis of Crowd Dynamics in the *Black Lives Matter* Online Social Movement

Hao Peng
University of Michigan
haopeng@umich.edu

Ceren Budak
University of Michigan
cbudak@umich.edu

Daniel M. Romero
University of Michigan
drom@umich.edu

ABSTRACT

Online social movements (OSMs) play a key role in promoting democracy in modern society. Most online activism is largely driven by critical offline events. Among many studies investigating collective behavior in OSMs, few has explored the interaction between crowd dynamics and their offline context. Here, focusing on the *Black Lives Matter* OSM and utilizing an event-driven approach on a dataset of 36 million tweets and thousands of offline events, we study how different types of offline events—police violence and heightened protests—influence crowd behavior over time. We find that police violence events and protests play important roles in the recruitment process. Moreover, by analyzing the re-participation dynamics and patterns of social interactions, we find that, in the long term, users who joined the movement during police violence events and protests show significantly more commitment than those who joined during other times. However, users recruited during other times are more committed to the movement than the other two groups in the short term. Furthermore, we observe that social ties formed during police violence events are more likely to be sustained over time than those formed during other times. Contrarily, ties formed during protests are the least likely to be maintained. Altogether, our results shed light on the impact of bursting events on the recruitment, retention, and communication patterns of collective behavior in the *Black Lives Matter* OSM.

KEYWORDS

online social movement, social media, black lives matter

ACM Reference Format:

Hao Peng, Ceren Budak, and Daniel M. Romero. 2019. Event-Driven Analysis of Crowd Dynamics in the *Black Lives Matter* Online Social Movement. In *Proceedings of the 2019 World Wide Web Conference (WWW '19)*, May 13–17, 2019, San Francisco, CA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3308558.3313673>

1 INTRODUCTION

Social movements often signify social change [5, 30, 35]. They help transform what have been considered radical ideas of one generation to the common sense of the next [14]. Social media has become instrumental in driving and facilitating politically contentious online social movements (OSMs) throughout the world [12, 20, 36, 44].

Previous research studied various OSMs such as the 2010 Arab Spring [28], the 2011 Egyptian Revolution [41, 50], the 2011 Occupy Wall Street [9, 20], through the lens of social media, mainly Twitter [2, 3, 22, 32, 40]. Topics they investigated range from participant motivation [44], movement framing [24, 43], informal learning [20], to information diffusion [22, 28, 41]. There are also studies focusing on crowd behavior in OSMs. Subjects they examined include recruitment process [23], user engagement [6, 9, 40], and online interactions [34], as well as the temporal evolution of crowd behavior [2, 4, 9, 11, 48]. Although these studies focus on different OSMs that were driven by very different political agendas, most of them convey the same message: *the unique characteristics of crowd behavior are often associated with critical offline events*. Yet, there is limited research on investigating how different types of offline events influence the behavior of online crowds and affect their interactions with the movement over time.

An event helps a person think about how the world diverges from his basic values [25]. On one hand, exogenous shocking events such as police shootings can contribute to the recruitment into the *Black Lives Matter* movement by creating “moral shocks” that can transform the public emotion from dread to outrage [17, 18, 25]. These complex emotional processes usually generate broader impact on the public sphere and can suddenly cause wide-scale unplanned movement participation of peripheral users in online social networks [8, 27, 31]. On the other hand, endogenous events such as self-organized protests can have an effect on crowd behavior with the help of movement networks in two stages. First, planned protests in physical locations are likely to be preceded by considerable communication and coordination among highly involved activists. Second, participants in offline protests often interact with the movement on social media to gain awareness, articulate their core beliefs, and offer a frame to the public [24]. As a result, offline events can generate massive participation in OSMs and are thus crucial to the emergence, growth, and decline of online activism [13].

In this paper, we analyze crowd dynamics in the *Black Lives Matter* (BLM) online movement using a combination of 36 million tweets and thousands of compiled offline events. We focus our analysis on crowd dynamics around three types of time periods: cases of police violence, large-scale offline protests, and periods of neither police violence or protests referred to as “other times”. Our event-driven investigation provides two major contributions. First, we find that police violence events and protests not only act as magnets to mobilize people around the movement, but they also greatly encourage user commitment in the long term. Second, we reveal that police violence events and protests have significantly different impact on the maintenance of social interactions established during these events. Altogether, our results shed light on how online crowds behave during bursting events.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '19, May 13–17, 2019, San Francisco, CA, USA

© 2019 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-6674-8/19/05.

<https://doi.org/10.1145/3308558.3313673>

2 RELATED WORK

Online Social Movements, Crowds, and Events. Social media is now used as an organizing tool for collective action and has enabled the proliferation of online low-risk activism [31], which may indeed have been a productive component of recent political uprisings [41]. In addition, studies have revealed that social media largely promotes social movements by means of influencing individual decisions about participation [20], providing new sources of information beyond the control of authorities, and shaping the discourse, logistics, and outcomes of protests [44]. This suggests that social media facilitates the emergence of massive offline mobilization and can serve as a good proxy of studying OSMs.

Several studies in OSMs have been conducted through the lens of Twitter [2, 6, 22, 28, 32, 41, 48]. Existing work tends to focus on the characteristics of information generated during the course of various online movements. They are interested in analyzing movements along dimensions including the variation of the amount of information produced, the way how information spreads [22, 28, 41], and the contribution from different social actors involved in the movement [2, 6, 32, 48].

Other research examines crowd behavior and its dynamics in OSMs. Topics covered range from online social interactions [34], the recruitment dynamics [23], and community polarization [50], to opinion evolution [4]. There are also studies exploring the connections between the dynamics of OSMs and relevant offline events [7, 52]. In [42], Steinert et al. studied the relationship between online activities and the volume of protests on the ground. The work in [3] explored the possibility of forecasting onsite protests using social media stream data. Spiro et al. found that key time points with unique online social structure often map to exogenous events in physical locations [40]. These studies inform us about the relationship between the movement dynamics and relevant protests unfolded on the ground. However, to our knowledge, there is little work investigating how different types of offline events affect crowd dynamics in OSMs.

Outside of movement research, the interplay between crowd behavior and shocking events has been a rising theme of considerable interest in setting such as organizations [37], crowdsourcing platforms [53, 54], and global scale mobile communications [1]. Bagrow et al. analyzed the behavioral change and societal response of crowds to external perturbation during emergencies [1]. Romero et al. investigated the link between stock market shocks and the structure of the communication network in a large hedge fund [37]. A recent work [53] studied the collective response of editors of Chinese Wikipedia to the government censorship in mainland China. These studies reveal that online crowds are especially sensitive during external perturbations [19, 21, 46]. The above literature thus forms a significant part of the motivation for our effort on analyzing crowd dynamics in OSMs using an event-driven approach.

Research in the BLM Movement. *BLM* is an international activist movement initiated by the African American community in the U.S. after the killing of Trayvon Martin in 2012. The movement is centered around issues related to racial inequality and police brutality. Since the extensive adoption of the hashtag *#BlackLivesMatter* on social media after the death of Eric Garner and Michael Brown

in 2014, the movement has generated thousands of on-site protests and demonstrations globally. Over time, the *BLM* movement gained worldwide recognition and the general public began to participate in the movement using a variety of hashtags on Twitter [24].

Several studies have analyzed the *BLM* online movement. In [33], Olteanu et al. provided a demographic characterization of users involved in this movement. Stewart et al. examined the frame process between competing online communities in the *BLM* movement through the use of related hashtags [43]. Ince et al. investigated how Twitter users interact with the *BLM* movement through hashtags and thus modify its framing [24]. A recent work [16] studied the divergent discourse between the *#BlackLivesMatter* protest and the *#AllLivesMatter* counter-protest on Twitter. Twyman et al. studied the dynamics of attention and knowledge production related to the *BLM* movement using Wikipedia [45]. However, as far as we know, there is limited work focusing on crowd dynamics in the *BLM* online movement and its relationship to offline events.

Most related to our work, De Choudhury et al. investigated the participation in the *BLM* online movement over time in different geographical locations and explored its relationship with protests unfolded on the ground [11]. Our work differs in that we systematically investigate the dynamics of user participation and patterns of social interactions, and explore how different types of offline events affect these dynamics in the *BLM* online movement.

3 DATA AND BACKGROUND

Twitter Dataset. Research indicates that the public use hashtags to interact with *BLM* to promote their messages, articulate core beliefs, and offer political solutions [11, 24, 33]. Here, we use data collected by [6] which includes tweets that are tagged with 53 *BLM* related hashtags. These race-related hashtags cover the stream of demonstrations and riots against the systemic racism and police violence toward African Americans [6].

The dataset spans from Jan. 2014 to May 2015, which covers the initial periods of the *BLM* movement on Twitter since it did not gain enough momentum until 2014 [11, 33]. It includes *every tweet* posted during this time period that contained one of the 53 *BLM* related hashtags mentioned above. There are 36 million tweets posted by 4 million users, among which we obtain 27 million retweets.

Offline Events. We compiled a list of *BLM* related events including police violence against African Americans and numerous on-site protests. The police violence events are based on news media reports and here we focus on the most prominent ones, which are shown in Table 1. The protest events are collected from *Elephrame*¹ [11]. This website keeps track of the *BLM* related demonstrations using information from multiple online sources. It estimates the number of demonstrators based on photos, videos, and media reports about each demonstration. Here we focus on protests with significant public awareness and potential impact on online crowds by defining a threshold $\mathcal{T}_{\text{protest}}$ of the number of participants and filtering those below it. We are able to identify 37 heightened protests with $\mathcal{T}_{\text{protest}} = 500$ in this study (The qualitative result does not change for $\mathcal{T}_{\text{protest}} = 100$).

¹<https://elephrame.com/textbook/BLM/chart>

Table 1: Five police violence events used in our analysis and their number of NYTimes articles.

Victim	Date	Location	# articles
Eric Garner	2014-07-17	New York City	518
Michael Brown	2014-08-09	Ferguson, Missouri	820
Tamir Rice	2014-11-23	Cleveland, Ohio	67
Walter Scott	2015-04-04	Charleston, South Carolina	47
Freddie Gray	2015-04-12	Baltimore, Maryland	169

Event Categorization. To explore the interplay between crowd behavior on Twitter and identified offline events in the *BLM* movement, we create an event-type categorization of the Twitter dataset.

For each instance of police violence, we label the week after the time when the event occurs as a *police violence* period. For each protest, we label the week before and the week after the event as a *protest* period, reflecting the possibility that participants use social media to organize and recruit for the upcoming protest [16]. We combine adjacent time periods of the same event type into a single period. When a protest period overlaps with a police violence period, it is likely that the former is immediately driven by the latter, thus the overlapping period is labeled as *police violence*. Time periods between different events are labeled as *other times*, which can be considered as normal times when public interactions with the movement are not induced by external factors. Throughout this study, we exclude tweets posted before July 17, 2014 since there is no protest data in *Elephrame* before this date. The event categorization is shown as the *background* in Figure 1.




Figure 1: Number of new/all participants in each of the non-overlapping, one-week time windows starting from the first police violence event in the offline dataset. (Background): Event categorization of the twitter dataset used in this analysis. Five dashed vertical lines represent five police violence events in Table 1. A rectangular area denotes a time period, and its color represents its event type.

We first analyze the recruitment process in the *BLM* online movement using this event categorization. We split the whole periods into non-overlapping bins of one week starting from the time of the police violence against Eric Garner. We define the arrival time D_u of user u , as the time when u tweeted for the first time. For a

given time window w_i , let n_i denote the number of new adopters whose arrival time D_u is during week w_i . Let p_i denote the number of participants who posted at least one tweet in week w_i . Figure 1 illustrates the time series of n_i and p_i respectively. It shows that there are repeating temporal variations in terms of the participation volume in the movement, with police violence events generating initial spikes followed by slight declines during protests and then fluctuating in other times. We also see that police violence events and protests function as magnets to attract new participants and mobilize people around the movement over time, which highlights the importance of examining mobilization patterns from the perspective of such events. The increase in activity during these events naturally leads to the question of whether users who became activated during different events behave and interact differently with the movement in the future.

4 DYNAMICS OF USER PARTICIPATION

Identifying characteristics of continued participation has been the central theme in recent OSM research [2, 6, 32, 48]. Yet, little is known about the evolution of user commitment as a function of offline events. In this section, we explore how arriving at the movement during a type of event is related to continued user participation in the *BLM* online movement.

We define a tracking window of t weeks to measure, for each user u , their commitment using their re-participation intensity within t weeks starting from D_u (their join time) along three dimensions. The first dimension is a binary variable y_u indicating if u re-participated or not within t weeks. The second one is the number of times f_u that u re-participated within t weeks (tweets in the movement). The third one is the number of days p_u that u re-participated within t weeks (active days in the movement). These three dimensions are meant to construct a thorough characterization of user commitment.

Modeling re-participation dynamics is complex. In addition to the events happening when they join, many other factors can directly influence users' re-participation behavior. To test whether the event during which users join the movement is related to their future participation, we utilize a prediction framework to control for potential confounds. For the first dimension, we adopt a logistic regression model where y_u is used as the dependent variable (Eq. 1). We employ a linear regression model for the second and the third dimensions separately, where $\log(f_u + 1)$ and $\log(p_u + 1)$ are used as the dependent variable, respectively (Eq. 2).

$$\text{logit}(p(y = 1)) = \beta_1 I + \sum_{k=2}^K \beta_k x_k + \epsilon \quad (1)$$

$$\log(y + 1) = \beta_1 I + \sum_{k=2}^K \beta_k x_k + \epsilon \quad (2)$$

The 8 independent variables are the same in all regressions, which are listed in Table 2. Variable 1 is categorical for which we use *other times* as the baseline, with two dummy variables indicating *police violence* and *protest*, respectively (I in Eq. 1 and Eq. 2). Variables 2 – 8 are numerical and are used as controls in the regressions. Variables 2 – 4 take into account the fact that users arrived at different time and the events following a user's start date can impact their re-participation dynamics irrespective of the events that occurred when they arrived. For instance, users may be more likely

Table 2: Independent variables for predicting user re-participation.

1	the event type of D_u (categorical)
2	the week number of D_u
3	the fraction of days labeled as <i>violence</i> in t weeks starting from D_u
4	the fraction of days labeled as <i>protest</i> in t weeks starting from D_u
5	the tweet count of u at D_u
6	the friend count of u at D_u
7	the follower count of u at D_u
8	the account age (in days) of u at D_u
Note: Variables 5 – 8 are user meta data.	

to return to the movement if there is an upcoming police violence event since these events usually encourage participation. Variables 5 – 8 control for the effect that can be explained by the difference between various user attributes. The Variance Inflation Factor (VIF) of all controls are below 5 in a multicollinearity test, which suggests that all of them should be included in the model. We standardize all numerical variables before performing the regression.

We train 10 models for each dimension by varying the tracking window from $t = 1$ to $t = 10$. To avoid right censoring bias, we focus on the same set of users whose join time D_u allow us to track their future participation for at least 10 weeks in our dataset. In total, we have 3.8 million users, among which 54% joined the BLM online movement during police violence periods, 42% joined during protest periods, and 4% joined during other times.

Our main objective is to estimate the impact of offline events on user commitment, therefore we focus on the first independent variable—the event type of users’ starting time, and measure the relative change for *police violence* and *protest* over *other times*. Namely, for each dummy variable, we calculate the percentage increase in each dimension over the baseline. In figure 2, we show the effects of events on user commitment in the BLM online movement.

How likely are users to re-participate? The results indicate that both police violence events and protests have a time-dependent effect on users’ likelihood of re-participation. We find that, in the long term, users who joined the movement during police violence and protest periods are significantly more likely to come back than those who joined during other times. Figure 2A shows that, compared to users who came during other times, the odds of re-participating in the movement within 8 – 10 weeks are nearly doubled for users who started in violence periods ($p < 10^{-10}$), while the odds of re-participation for users who joined during protests is 50% higher than the baseline group (other times) ($p < 10^{-10}$).

However, in the short term, the odds of re-participating in the movement decrease by about 40% for users in the protest group over the baseline ($p < 10^{-10}$), and users in the other two groups are comparable overall with no clear dominance of one group over the other in terms of the likelihood of re-participation.

How frequently and how persistently do they re-participate? Merely assessing the presence or absence of continued participation in an activist movement is not enough to measure user dedication. We thus further examine users’ commitment by investigating their participation frequency and persistence, namely the number of




Figure 2: (A) The percentage increase in odds of re-participation within t weeks for users who joined during police violence and protest periods over those who joined during other times. The error bars indicate a 95% confidence interval. (B) The same as (A), but for the percentage increase in users’ re-participation times (Frequency dimension). (C) The same as (A), but for the percentage increase in users’ re-participation days (Persistence dimension).

tweets a user has posted and the number of days that he has participated in a given time window. The latter complements the former by capturing more nuanced participation dynamics—some users can post up to hundreds of tweets a day but stay silent thereafter.

We find that police violence events and protests have similar two-fold effects on users’ future commitment. According to Figure 2 B and Figure 2C, in the long term, joining the movement during violence periods contributes to an approximate 60% increase in re-participation frequency and a 40% increase in persistence over those recruited during other times, and protest group experiences about a 25% and a 20% increase, respectively. In the short term, however, the two groups have a 10% – 20% decrease in the frequency and the persistence dimensions compared to other times group.

Discussion of findings. Overall, we find that police violence events and heightened protests significantly affect individual participation in the BLM online movement. These events have both short and long term effects on the dynamics of user commitment. One possible explanation is that users who joined during other times are

self-motivated and thus are more active in the initial periods, but their energy decreases over time. On the other hand, arriving to the movement during high activity events, when the entire movement is more energized, may influence users to be less active in the short term as they process the large volume of information they are being exposed to, but more active in the long term. Indeed, studies have shown that attending protests leads to higher long term political and protest participation [29, 39]. Another plausible reason is that people who joined during different periods may respond differently to changes in the movement composition over time [26, 51]. Alternatively, people joining the movement during different types of events can have different motivations and thus different responses to insufficient gratification that affect long term movement commitment [47]. We acknowledge that future work is needed to determine which, if any, of these factors drive the time-dependent participation response.

5 PATTERNS OF SOCIAL INTERACTIONS

Individual participation is only one component of OSMs. The interactions among movement participants make the community strong and robust [19]. It has been suggested that crowds display higher strong tie interactions and interpersonal exchanges during shocks and societal upheavals [19, 37]. But to what extent do external events affect tie strength over time? In this section, we characterize the formation and maintenance of social ties formed during different types of offline events in the *BLM* online movement.

Specifically, we consider retweets as social interactions between participants involved in the movement. For a social tie $e_{i,j} = (u_i, u_j)$, we define $D_{e_{i,j}}$ as the first time u_i retweeted a tweet by u_j that contains one of the hashtags. We measure tie strength using its interaction intensity within t weeks (the tracking window) starting from $D_{e_{i,j}}$ along three dimensions similar to the ones we applied to user re-participation. The first dimension is a binary variable $y_{e_{i,j}}$ indicating if u_i retweeted u_j again within t weeks ($e_{i,j}$ recurred or not). The second one is the number of times $f_{e_{i,j}}$ that u_i retweeted u_j again within t weeks (the recur frequency of $e_{i,j}$). The third one is the number of days $p_{e_{i,j}}$ that u_i retweeted u_j again within t weeks (the recur persistence of $e_{i,j}$). These three dimensions can provide a rich description of the strength of social interactions.

Similar to user participation, many factors can influence the patterns of social interactions. To test whether the creation time (event type) of a tie is related to its recur patterns, we leverage a prediction framework to control for possible confounding factors. We follow the same pipeline for each dimension as in Section 4 (Eq. 1 for the first dimension and Eq. 2 for the other two).

The 20 independent variables are the same in all regressions, which are listed in Table 3. We define D_0 as the time of the first tweet in our dataset. Variables 2–20 are used as controls in the regressions, which capture the variations between ties’ creation time in the movement (4, 19, 20), the tendency for users to retweet (11, 13), the tendency for users to be retweeted (12, 14), users’ activity level prior to and after tie creation (5 – 10, 17, 18), number of users they retweeted in common (16), and users previous interactions (15), etc.

The Variance Inflation Factor (VIF) of all numerical variables (4 – 20) are below 5 in a multicollinearity test. We standardize all numerical variables before performing the regression. Again, we

Table 3: Independent variables for predicting tie strength.

1	the event type of $D_{e_{i,j}}$
2	the event type of D_{u_i}
3	the event type of D_{u_j}
4	the week number of $D_{e_{i,j}}$
5	the # protest tweets of u_i in $[D_0, D_{e_{i,j}}]$
6	the fraction of (5) in <i>violence</i>
7	the fraction of (5) in <i>protest</i>
8	the # protest tweets of u_j in $[D_0, D_{e_{i,j}}]$
9	the fraction of (8) in <i>violence</i>
10	the fraction of (8) in <i>protest</i>
11	the times u_i has retweeted others in $[D_0, D_{e_{i,j}}]$
12	the times u_i has been retweeted by others in $[D_0, D_{e_{i,j}}]$
13	the times u_j has retweeted others in $[D_0, D_{e_{i,j}}]$
14	the times u_j has been retweeted by others in $[D_0, D_{e_{i,j}}]$
15	the times u_j has retweeted u_i in $[D_0, D_{e_{i,j}}]$
16	the # users u_i and u_j have retweeted in common in $[D_0, D_{e_{i,j}}]$
17	the # protest tweets of u_i in t weeks from $D_{e_{i,j}}$
18	the # protest tweets of u_j in t weeks from $D_{e_{i,j}}$
19	the fraction of days labeled as <i>violence</i> in t weeks from $D_{e_{i,j}}$
20	the fraction of days labeled as <i>protest</i> in t weeks from $D_{e_{i,j}}$
Note: Variables 5 – 20 are calculated based on our dataset.	

train 10 models for each dimension by varying the tracking window from $t = 1$ to $t = 10$. We focus on the same set of social ties whose creation time $D_{e_{i,j}}$ allow us to track their future interactions for at least 10 weeks in the dataset. In total, we obtain 17 million social ties, among which 41% were created during police violence events, 53% during protest periods, and 6% in other times.

We apply the same procedure to the first independent variable as we did in Section 4. We measure the relative change in the strength for social ties formed during police violence events and protests over those in the baseline group. The results are shown in Figure 3.

How likely are interactions to recur? We find that social interactions started in police violence periods are significantly more likely to be sustained in the long term than those created during other times. Figure 3A shows that, within 8 – 10 weeks, the odds of recur for social ties created in violence periods are 10% – 20% higher than that of other times group ($p < 10^{-15}$). However, the short term differences between two groups are less pronounced as there is oscillation between them. On the other hand, protests have a more unified influence on the maintenance of social interactions, i.e., irrespective of the time effect, communication bonds established in protests are less likely to recur than do those started in other times (about 10% decrease overall, $p < 10^{-9}$).

How frequently and how persistently do they recur? We further examine tie strength along the frequency and persistence dimensions. Figure 3B and 3C show that the overall influence of police violence events and protests on the recur frequency and persistence of social ties are similar to what we observe in the first dimension, but the result also indicates that violence events and protests have marginal influence on the sustainability of social ties compared to the baseline group (only 1% increase or decrease, $p < 10^{-10}$). This may be explained by the fact that, in general, people communicate




Figure 3: (A) The percentage increase in odds of recur for social interactions started during police violence and protest events over those created during other times. The error bars indicate a 95% confidence interval. (B) The same as (A), but for the percentage increase in recur frequency. (C) The same as (A), but for the percentage increase in recur persistence.

less selectively and form short-lived connections during political upheavals [38], thus reducing the influence of external events on tie strength and its evolution.

Discussion of findings. In summary, we find that communications started in police violence events are more likely to be maintained over time than interactions formed in normal times, which supports the idea that communities bond and build strong ties during shocks [15, 19, 37]. However, social ties formed during protests recur less likely, less frequently and less persistently than do those established in other times. We conjecture that social connections forged during protests are more about informational support which may be less likely to persist compared to emotional support [10, 49].

6 CONCLUSION

We conduct a large-scale empirical investigation of crowd dynamics in the *BLM* online movement using a combination of social media data and a comprehensive list of offline events. By analyzing

user participation and social interactions through the lens of different types of events, we are able to reveal the dynamics of online crowd behavior and its relationship to offline events in a politically contentious movement around race.

First, we find that both police violence events and intensified protests have time-dependent effects on users' future commitment. In the long term, users who joined the movement during two type of events are significantly more likely to re-participate than those who joined during other times. However, in the short term, joining the movement in these events are associated with decreased future commitment. **Second**, we find that police violence events and protests have different effects on the continuation of social interactions formed during these periods, with the former showing both long-term and short-term effects while the latter displaying a more unified pattern. Specifically, we find that communications started in police violence events are significantly more likely to be sustained in the long run than those started in other times. But in the short term, they are less likely to be maintained. However, regardless of the time, social ties established during protests are less likely to recur compared to other times group.

Implications. Our study has implications for policymakers, movement organizers, and online social movement observers. These escalated events, unexpected or expected, function as driving forces that organize people to participate in the movement over time. They also play a great role in user retention and motivating social interactions. Knowing about these patterns, movement organizers and organizations may allocate more resources on the recruitment process during normal times. A social movement similar to the *BLM* would also try to encourage newcomers who arrived when no special event was happening to sustain participation in the long term, since they are known to drop out over time. Similarly, cultivating stronger social relationships formed during protests may be an effective way to stimulate sustained communications and facilitate information sharing about the movement in the future. Finally, this study provides insights to research on collective behavior and crowd response in the face of bursting events.

Limitations and future work. This work has limitations. First, we measure user commitment solely based on twitter dataset. However, it's unclear whether individuals who participated in the movement at its initial stage did not divert their participation to other places that are not captured in our data. Second, users may participate on Twitter in ways other than using hashtags. Third, we only study one online social movement, which limits the generalizability of our findings to other type of movements. Furthermore, although we find strong correlation between online crowd dynamics and different types of offline events, we cannot make causal claims based on these findings. Investigating whether users were intrinsically different or they were just influenced by the events when they joined the movement is an open direction for future research.

ACKNOWLEDGMENTS

This research was partly supported by the National Science Foundation Grants No. IIS-1617820 and IIS-1815875. We would like to thank David Jurgens for helpful discussion, and Lia Bozarth for data preparation.

REFERENCES

- [1] James P Bagrow, Dashun Wang, and Albert-Laszlo Barabasi. 2011. Collective response of human populations to large-scale emergencies. *PLOS ONE* 6, 3 (2011).
- [2] Pablo Barberá, Ning Wang, Richard Bonneau, John T Jost, Jonathan Nagler, Joshua Tucker, and Sandra González-Bailón. 2015. The critical periphery in the growth of social protests. *PLOS ONE* 10, 11 (2015).
- [3] Marco T Bastos, Dan Mercea, and Arthur Charpentier. 2015. Tents, tweets, and events: The interplay between ongoing protests and social media. *Journal of Communication* 65, 2 (2015), 320–350.
- [4] Javier Borge-Holthoefer, Walid Magdy, Kareem Darwish, and Ingmar Weber. 2015. Content and network dynamics behind Egyptian political polarization on Twitter. In *CSCW*. ACM.
- [5] Philip J Bostic. 2013. Social Movement. (2013). <https://www.learningtogive.org/resources/social-movement>
- [6] Lia Bozarth and Ceren Budak. 2017. Is Slacktivism Underrated? Measuring the Value of Slacktivists for Online Social Movements. In *ICWSM*.
- [7] Ceren Budak and Duncan J Watts. 2015. Dissecting the spirit of Gezi: Influence vs. selection in the Occupy Gezi movement. *Sociological Science* 2 (2015), 370–397.
- [8] Michael Conover, Jacob Ratkiewicz, Matthew R Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on twitter. *ICWSM* (2011).
- [9] Michael D Conover, Emilio Ferrara, Filippo Menczer, and Alessandro Flammini. 2013. The digital evolution of occupy wall street. *PLOS ONE* 8, 5 (2013).
- [10] Carolyn E Cutrona and Julie A Suhr. 1992. Controllability of stressful events and satisfaction with spouse support behaviors. *Communication Research* 19, 2 (1992), 154–174.
- [11] Munmun De Choudhury, Shagun Jhaver, Benjamin Sugar, and Ingmar Weber. 2016. Social Media Participation in an Activist Movement for Racial Equality. In *ICWSM*.
- [12] Donatella Della Porta and Alice Mattoni. 1999. *Social movements*. Wiley Online Library.
- [13] Mario Diani. 2003. *Social movements and networks*. Wiley Online Library.
- [14] Peter Dreier. 2012. Social Movements: How People Make History. *Mobilizing Ideas* (2012). <https://mobilizingideas.wordpress.com/2012/08/01/social-movements-how-people-make-history>
- [15] Kimberly Falconer, Mieke Sachsenweger, Kerry Gibson, and Helen Norman. 2011. Grieving in the internet age. *New Zealand Journal of Psychology* 40, 3 (2011).
- [16] Ryan J Gallagher, Andrew J Reagan, Christopher M Danforth, and Peter Sheridan Dodds. 2018. Divergent discourse between protests and counter-protests: #BlackLivesMatter and #AllLivesMatter. *PLOS ONE* 13, 4 (2018).
- [17] William A Gamson. 1992. *Talking politics*. Cambridge university press.
- [18] William A Gamson, Bruce Fireman, and Steven Rytina. 1982. *Encounters with unjust authority*. Dorsey Press.
- [19] Kimberly Glasgow, Clayton Fink, and Jordan L Boyd-Graber. 2014. "Our Grief is Unspeakable": Automatically Measuring the Community Impact of a Tragedy. In *ICWSM*.
- [20] Benjamin Gleason. 2013. # Occupy Wall Street: Exploring informal learning about a social movement on Twitter. *American Behavioral Scientist* 57, 7 (2013), 966–982.
- [21] Lewis R Goldfrank, Allison M Panzer, Adrienne Stith Butler, et al. 2003. *Preparing for the psychological consequences of terrorism: A public health strategy*. National Academies Press.
- [22] Sandra González-Bailón, Javier Borge-Holthoefer, and Yamir Moreno. 2013. Broadcasters and hidden influentials in online protest diffusion. *American Behavioral Scientist* 57, 7 (2013), 943–965.
- [23] Sandra González-Bailón, Javier Borge-Holthoefer, Alejandro Rivero, and Yamir Moreno. 2011. The dynamics of protest recruitment through an online network. *Scientific Reports* 1 (2011).
- [24] Jelani Ince, Fabio Rojas, and Clayton A Davis. 2017. The social media response to Black Lives Matter: how Twitter users interact with Black Lives Matter through hashtag use. *Ethnic and Racial Studies* 40, 11 (2017), 1814–1830.
- [25] James M Jasper. 1998. The emotions of protest: Affective and reactive emotions in and around social movements. *Sociological Forum* 13, 3 (1998), 397–424.
- [26] Bert Klandermans. 1994. Transient identities? Membership patterns in the Dutch peace movement. *New social movements: From ideology to identity* (1994), 168–184.
- [27] Fannie Liu, Denae Ford, Chris Parnin, and Laura A. Dabbish. 2017. Selfies as Social Movements: Influences on Participation and Perceived Impact on Stereotypes. *PACMHCI* 1 (2017).
- [28] Gilad Lotan, Erhardt Graeff, Mike Ananny, Devin Gaffney, Ian Pearce, et al. 2011. The Arab Spring| the revolutions were tweeted: Information flows during the 2011 Tunisia and Egyptian revolutions. *International Journal of Communication* 5 (2011), 31.
- [29] Doug McAdam. 1989. The biographical consequences of activism. *American sociological review* (1989), 744–760.
- [30] Doug McAdam, John D McCarthy, and Mayer N Zald. 1996. *Comparative perspectives on social movements: Political opportunities, mobilizing structures, and cultural framings*. Cambridge University Press.
- [31] Evgeny Morozov. 2009. The brave new world of slacktivism. *Foreign Policy* 19, 05 (2009).
- [32] Christine Ogan and Onur Varol. 2017. What is gained and what is left to be done when content analysis is added to network analysis in the study of a social movement: Twitter use during Gezi Park. *Information, Communication & Society* 20, 8 (2017), 1220–1238.
- [33] Alexandra Olteanu, Ingmar Weber, and Daniel Gatica-Perez. 2016. Characterizing the demographics behind the #blacklivesmatter movement. In *2016 AAAI Spring Symposium Series*.
- [34] Leysia Palen and Sarah Vieweg. 2008. The emergence of online widescale interaction in unexpected events: assistance, alliance & retreat. In *CSCW*. ACM.
- [35] Pamela Paxton. 2002. Social capital and democracy: An interdependent relationship. *American Sociological Review* (2002), 254–277.
- [36] Tom Postmes and Suzanne Brunsting. 2002. Collective action in the age of the Internet: Mass communication and online mobilization. *Social Science Computer Review* 20, 3 (2002), 290–301.
- [37] Daniel M Romero, Brian Uzzi, and Jon Kleinberg. 2016. Social networks under stress. In *WWW*.
- [38] Cuihua Shen, Peter Monge, and Dmitri Williams. 2014. The evolution of social ties online: A longitudinal study in a massively multiplayer online game. *Journal of the Association for Information Science and Technology* 65, 10 (2014), 2127–2137.
- [39] Darren E Sherkat and T Jean Blocker. 1997. Explaining the political and personal consequences of protest. *Social Forces* 75, 3 (1997), 1049–1070.
- [40] Emma S Spiro and Andrés Monroy-Hernández. 2016. Shifting Stakes: Understanding the Dynamic Roles of Individuals and Organizations in Social Media Protests. *PLOS ONE* 11, 10 (2016).
- [41] Kate Starbird and Leysia Palen. 2012. (How) will the revolution be retweeted?: information diffusion and the 2011 Egyptian uprising. In *CSCW*. ACM.
- [42] Zachary C Steinert-Threlkeld, Delia Mocanu, Alessandro Vespignani, and James Fowler. 2015. Online social networks and offline protest. *EPJ Data Science* 4, 1 (2015).
- [43] Leog Stewart, Ahmer Arif, A Conrad Nied, Emma S Spiro, and Kate Starbird. 2017. Drawing the Lines of Contention: Networked Frame Contests Within #BlackLivesMatter Discourse. *Proc. ACM Human Computer Interaction* 1 (2017).
- [44] Zeynep Tufekci and Christopher Wilson. 2012. Social media and the decision to participate in political protest: Observations from Tahrir Square. *Journal of Communication* 62, 2 (2012), 363–379.
- [45] Marlon Twyman, Brian C Keegan, and Aaron Shaw. 2017. Black Lives Matter in Wikipedia: Collective memory and collaboration around online social movements. In *CSCW*.
- [46] Robert J Ursano, Carol S Fullerton, Tzu-Cheg Kao, and Vivek Bhartiya. 1995. Longitudinal assessment of posttraumatic stress disorder and depression after exposure to traumatic death. *Journal of Nervous and Mental Disease* (1995).
- [47] Jacquelin Van Stekelenburg and Bert Klandermans. 2017. Individuals in movements: A social psychology of contention. In *Handbook of social movements across disciplines*. Springer, 103–139.
- [48] Onur Varol, Emilio Ferrara, Christine L Ogan, Filippo Menczer, and Alessandro Flammini. 2014. Evolution of online user behavior during a social upheaval. In *WebSci*. ACM.
- [49] Zijian Wang and David Jurgens. 2018. It's going to be okay: Measuring Access to Support in Online Communities. In *EMNLP*.
- [50] Ingmar Weber, Venkata R Kiran Garimella, and Alaa Batayneh. 2013. Secular vs. islamist polarization in egypt on twitter. In *Proceedings of the ACM International Conference on Advances in Social Networks Analysis and Mining*. ACM, 290–297.
- [51] Nancy Whittier. 1997. Political generations, micro-cohorts, and the transformation of social movements. *American Sociological Review* (1997), 760–778.
- [52] Donghyeon Won, Zachary C Steinert-Threlkeld, and Jungseok Joo. 2017. Protest Activity Detection and Perceived Violence Estimation from Social Media Images. In *Proceedings of the 2017 ACM on Multimedia Conference*. ACM.
- [53] Ark Fangzhou Zhang, Danielle Livneh, Ceren Budak, Lionel P. Robert, and Daniel M. Romero. 2017. Shocking the Crowd: The Effect of Censorship Shocks on Chinese Wikipedia. In *ICWSM*.
- [54] Ark Fangzhou Zhang, Danielle Livneh, Ceren Budak, Lionel P. Robert, Jr., and Daniel M. Romero. 2017. Crowd Development: The Interplay Between Crowd Evaluation and Collaborative Dynamics in Wikipedia. *PACMHCI* 1 (2017).